

AD-A112 936

STANFORD UNIV CA DEPT OF COMPUTER SCIENCE F/6 9/2
THE DEVELOPMENT OF REPRESENTATION METHODS TO FACILITATE KNOWLEDGE--ETC(U)
DEC 81 E H SHORTLIFFE, M R GENESERETH N00014-81-X-0004

UNCLASSIFIED

NL

END
DATE
FILMED
4 82
DTIC

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

12

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO. ND-4112936	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) See attached list	5. TYPE OF REPORT & PERIOD COVERED Technical - Annual progress report	
	6. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) See attached list	8. CONTRACT OR GRANT NUMBER(s) ONR N00014-81-K-0004	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Computer Science Stanford University Stanford, CA 94305		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 154-436
11. CONTROLLING OFFICE NAME AND ADDRESS Program Director Information Systems Mathematical & Information Sciences Div. ONR; 800 N. Quincy St., Arlington, VA		12. REPORT DATE December 1981
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) ONR Resident Representative Durand Aeronautics Bldg., Rm. 165 Stanford University Stanford, California 94305		13. NUMBER OF PAGES see attached list
16. DISTRIBUTION STATEMENT (of this Report) See attached		15. SECURITY CLASS. (of this report) Unclassified
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		

82 04 02 001

DTIC
ELECTE
APR 01 1982

E

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102- LF-014-6601UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

ADA112936

DTIC FILE COPY

December 31, 1981

Dr. Marvin Denicoff, Program Director
Information Sciences Program
Office of Naval Research
800 N. Quincy Street
Arlington, Virginia 22217



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	<i>per file</i>
By	
Distribution/	
Availability Codes	
Dist	Special
<i>A</i>	

ANNUAL PROGRESS REPORT

ONR TASK AUTHORITY IDENTIFICATION NUMBER: NR 049-479

CONTRACTOR: The Board of Trustees of the Leland Stanford Junior University

CONTRACT NO.: N00014-81-K-0004

CONTRACT AMOUNT: \$456,622

EFFECTIVE DATE OF CONTRACT: January 1, 1981

EXPIRATION DATE OF CONTRACT: December 31, 1983

PRINCIPAL INVESTIGATOR: Edward H. Shortliffe

TELEPHONE NO.: (415) 497-6979

SHORT TITLE OF WORK: The Development of Representation Methods to Facilitate
Knowledge Acquisition and Exposition in Expert Systems

REPORTING PERIOD: January 1, 1981 - December 31, 1981

* DESCRIPTION OF PROGRESS: See attached

* SCHEDULE FOR NEXT QUARTER: See attached

PROBLEMS ENCOUNTERED: none

DEPARTURES FROM PROPOSED WORK: none

ACTION REQUIRED BY THE GOVERNMENT: none

* FISCAL STATUS:

- (1) Amount currently provided on contract: \$140,825
- (2) Expenditures and commitments as of 12/31/81: \$113,001
- (3) Funds required to complete work: \$343,621

Stanford SPO Reference No.:
11192-01-00

Edward H. Shortliffe
Edward H. Shortliffe, MD, PhD
Principal Investigator

The Development of Representation Methods to Facilitate
Knowledge Acquisition and Exposition in Expert Systems

ONR Contract N00014-81-K-0004

Principal Investigator: Edward H. Shortliffe, M.D., Ph.D.

Associate Investigator: Michael R. Genesereth, Ph.D.

1 PROGRESS REPORT (January 1, 1981 - December 31, 1981)

During the first year of research, we have studied the design and development of optimal techniques for the computer-based representation of expert knowledge. As discussed in our proposal, we have concentrated on issues of knowledge acquisition and explanation. Because the underlying choice of a knowledge representation is key, much of the early work has involved experiments with alternate representation techniques. In this progress report we have organized our discussion around the individual subprojects that are underway. It should be noted, however, that there is a unifying research theme (knowledge representation for knowledge acquisition and explanation) that ties together the projects and the participating scientists and students. When technical reports have been prepared in the last year, we have indicated so in the discussion and included the reports as attachments to this summary.

1.1 MULTIPLE REPRESENTATION SYSTEM (MRS)

1.1.1 Knowledge Representation

The primary result of our research on representation this year has been the development of a representation language known as MRS. Two technical reports describing this system have been prepared and are included as attachments. MRS is a knowledge representation system for use by computer programmers who are building expert systems. The key feature of the system is its single powerful language for stating facts while storing them in a variety of different representations. Because of its multiple representations and the resulting flexible techniques for encoding knowledge, it has computational advantages over other knowledge representation systems. In addition, because of its flexible inference capability, we believe it offers an expressive superiority when compared to traditional data definition languages.

A second important feature of the system is its ability to represent and manipulate meta-level knowledge. In MRS the application program itself is treated as a domain about which knowledge can be encoded. One can write sentences about sentences and about the structure of the system; the

PROGRESS REPORT (January 1, 1981 - December 31, 1981)

inference system can then reason with these sentences just as it reasons about external domains like geology and medicine. In practice, MRS uses meta-level information to decide how to carry out each operation it is asked to perform. Thus, one can easily switch representations or inference methods by changing the appropriate sentences, and one can thereby enable the system to control its own operation.

MRS is implemented in a variety of Lisps, including MacLisp and Interlisp on the DEC-20, Interlisp-D on the Xerox Dolphin, and Franz-LISP on the Vax. The system is being used in a variety of projects at Stanford and elsewhere (e.g., automated computer diagnosis and intelligent operating system interface) and is now being taught in the first graduate level AI course at Stanford.

1.1.2 Problem Reformulation

An important MRS subproject is Mr. Steven Tappel's work on problem reformulation. Changing coordinate systems in solving three dimensional problems is a common example of the technique. Mr. Tappel is interested in automating this process. As part of that interest, he is trying to automate the selection of data structures in MRS. This involves reasoning about each representation to determine whether it preserves enough information to solve the desired class of problems. The task also requires strategies for interfacing different representations. Furthermore, it requires strategies for comparing and evaluating alternatives. The key determinant in choosing a good data representation is the capability of the underlying system. Tappel's work is especially timely as research in VLSI design makes available computers with widely varying capabilities.

1.1.3 Causality, Design, and Mechanistic Reasoning

We are also involved in an effort to apply MRS to the representation of knowledge about design and causality. The research on design description is being done in conjunction with an IBM supported project aimed at the automated diagnosis of computer hardware failures (the DART project). The primary result of this work has been the development of a design description language called SUBTLE and a coordinated reasoner/simulator called SHAM. SUBTLE enables its user to specify complete or partial information about a device's structure (its parts and their interconnections), its behavior, and its teleology (arguments showing how the structure gives rise to the behavior). SUBTLE uses MRS as its representational formalism but also includes a large number of terms useful in the description of circuit design. SHAM is essentially an "interpreter" for SUBTLE and can be used to explore the consequences of a design and thereby detect incompleteness or inconsistency. A manual for the SUBTLE language is attached to this report, and documentation on SHAM should be available early next year.

The work on causality is focussed on the representation of various aspects of human physiology. Mr. John Kunz is undertaking thesis research

PROGRESS REPORT (January 1, 1981 - December 31, 1981)

that uses SUBTLE to represent pulmonary and renal pathophysiology (see below). We have also begun looking at SUBTLE as an approach to representing the knowledge of endocrinology. It is our observation that there is considerable overlap between issues encountered in the diagnosis of computer hardware failures and causal reasoning to diagnose diseases or explain basic physiology. The early work suggests that it should be possible to apply the diagnostic techniques developed in the DART project directly to descriptions of human physiology. If so, it may be possible to build a medical diagnosis program that can articulate its reasoning at a deeper level than is possible in existing programs.

1.2 EXPERT EXPLANATIONS PROJECT (EXPEX)

The Expert Explanation project is a series of experiments, mostly in medical domains, that seek to deal with challenging issues in knowledge representation and the generation of explanations for users of expert systems. Areas of investigation include:

- 1) physiological modeling of renal and pulmonary function; these are clinical domains in which mathematical descriptions interact with symbolic knowledge to aid in predicting deviations from normal when reasoning from first principles.
- 2) clinical endocrinology, with an emphasis on calcium metabolism; this domain has been selected because the knowledge of endocrinology is unusually well-specified, with causal mechanisms generally well-delineated and feedback control leading to important questions of how best to model equilibrium states.
- 3) clinical oncology, particularly the modeling of expert decision making regarding optimal adjustment of cancer therapy depending upon the response to prior treatment; one challenge in this domain is to better understand how to represent temporal trends of pertinent parameters and, in turn, how to reason from an analysis of those trends.

1.2.1 Physiologic Modeling

Mr. John Kunz is working in the first of these areas. His goal is to develop a system which analyzes physiological situations using a model of underlying mechanisms. Such a system could identify potential causes of, changes to, and consequences of physiologically abnormal situations. The model is being constructed using the MRS system described briefly above. In addition to the model of the mechanism, his system specifies a process for analyzing that mechanism in order to identify its possible response to various situations that could arise in a patient. The combination of a model and an analysis process allows the program to identify the physiological rationale for diagnostic tests, medical diagnoses, alternative therapies, and proposed tests of the effectiveness of therapy.

PROGRESS REPORT (January 1, 1981 - December 31, 1981)

Mr. Kunz's system represents physiological functions -- such as mass transport and urine production -- and concepts of physics on which these functions depend, such as fluid flow in response to pressure gradients and conservation of mass. Modeling the physics of mechanisms allows consistent analysis of related situations, even though they may occur in different organ systems or clinical situations. In addition, by representing and manipulating basic physiological functions, we believe it will be possible to model higher-level systems more easily than it would be to model those functions separately. It should be stressed that the issues addressed in this work are not inherently medical; similar problems arise in the construction of expert systems in any domain in which detailed causal relationships are amenable to formal mathematical description coupled with symbolic models.

1.2.2 Explanations from Causal Relationships

We have also undertaken experiments in the generation of explanations using knowledge structures other than MRS. In one component of this research, largely the work of Dr. Jerry Wallis, we have focussed on theoretical issues regarding techniques for customizing explanations to the needs and expertise of the user of an expert system. Causal knowledge has been encoded in production rules that link conceptual nodes in a semantic network. Individual concepts are associated with static measures of their complexity and importance. When the system generates explanations from a reasoning chain of rules, it prunes the chain in accordance with the expertise of the user and knowledge about the complexity of the knowledge being discussed. A paper describing this work was submitted for publication in December; the corresponding technical report is attached.

1.2.3 Hypothesis Analysis for Decision Support

In the area of endocrinology, Mr. Greg Cooper has recently initiated work that seeks to develop methods for implementing expert systems that critique the ideas and opinions of their users rather than produce advice de novo. The work raises important issues regarding hypothetical reasoning.

Mr. Cooper's goal is to develop a system that receives both a medical hypothesis and patient data as input, then returns an analysis of the hypothesis in light of the data. It is the ability of the user to control the reasoning focus of the computer that largely distinguishes this kind of decision support system from most previous expert consultation systems. In most expert systems, the user is allowed to request a justification for what the program has done, but seldom is he allowed to redirect the system's problem solving focus.

The reason for focussing on this kind of decision support system is that it allows us to concentrate our efforts on hypothesis analysis, an issue that we feel is crucial for the development of high performance expert systems that are able to explain their reasoning. Since diagnosis is a

PROGRESS REPORT (January 1, 1981 - December 31, 1981)

critical component of medical decision making, and hypothesis analysis is similarly central to diagnosis, we feel that important gains can be made by concentrating our efforts on the more limited problem of hypothesis analysis. An analysis system that considers hypotheses only if they are generated by the machine can make certain assumptions about the nature and form of those hypotheses. However, an analysis system that can accept an arbitrary user hypothesis is able to make fewer such assumptions and thus must be more general.

Our philosophy dictates that the analyzer receive an hypothesis and patient data in a language that is comfortable to the user. It is equally important that the output of an analysis be in a form that is most meaningful to the user. The analysis is essentially an explanation of the plausibility of the user-posed hypothesis in light of current patient data.

The initial analysis system we have constructed is based on the knowledge base for the INTERNIST/CADUCEUS program. Although this knowledge base is very broad, much key knowledge is missing. We have therefore chosen to augment the endocrinology portion of the knowledge base -- in particular that portion dealing with calcium metabolism -- so that it is in a form that can be easily accessed for hypothesis analysis purposes. A simple decision support system has been developed which takes as input a disease and a list of patient manifestations, then outputs an analysis in terms of which manifestations are consistent, inconsistent, or unknown with respect to the disease hypothesis.

1.2.4 Psychological Studies of Explanation

Our interest in improving the explanation capabilities of expert systems has led us to fruitful collaboration and exchange with psychologists at Stanford who are interested in natural explanation among human beings. It is our belief that useful design criteria for expert systems can follow from formal psychologic experiments. One member of our group, Mr. Randy Teach, is undertaking a doctoral dissertation in educational psychology under partial support from the NIH. The project involves a study of the explanatory and reasoning patterns used by physicians. Six physician subjects are being asked to "think aloud" while solving 12 difficult diagnostic problems in endocrinology. They are then asked to dictate consultation letters back to an imagined "referring physician". These letters and the accompanying reasoning transcripts provide the basis for an analysis of the ways experts explain their reasoning to one another. The results of this work will be useful in guiding the design of explanation capabilities for future consultation programs to be used either by physicians or by other users of expert systems.

1.2.5 Temporal Reasoning and Explanation

Another major research project at Stanford, initiated under a grant from the National Library of Medicine, is the applications system known as

PROGRESS REPORT (January 1, 1981 - December 31, 1981)

ONCOCIN [1]. This large program is designed to be used by cancer experts who are managing patients undergoing treatment for tumors. Such patients are typically treated as outpatients, and the selection of therapy at the time of any given clinic visit is dependent upon an analysis of the patient's response to therapy on previous visits and a prediction of how he or she will respond to current therapy. A prototype system is already in use in the oncology clinic at Stanford Medical Center, but we are finding that ONCOCIN is also providing a challenging arena for the study of fundamental questions of interest to artificial intelligence researchers.

Two areas of particular relevance to our work on this contract are the problems of temporal reasoning and the associated explanation challenges that are inherent in the ONCOCIN domain. The problem of reasoning about parameters that vary over time is frequently encountered in non-medical domains as well, and we are hopeful that lessons learned in the context of ONCOCIN can be applied to problems such as the physiological modeling that John Kunz is undertaking. ONCOCIN is based on production rules similar to those used in EMYCIN systems [2], but the developing techniques for accessing a database of prior information and performing temporal trend analysis could be adapted for transfer to other representation schemes. ONCOCIN currently does not consider more than the two immediately preceding clinic visits in reaching a decision about current therapy; we intend to generalize this approach so that the patient's entire history is considered.

As ONCOCIN's reasoning based on prior therapy becomes more complex, we expect to encounter significant new challenges in generating explanations that are acceptable to the users of the system. Current explanations give justifications for advice on the basis of individual rules (in much the same way that EMYCIN systems do). This approach will be inadequate when trend analysis is introduced since we anticipate that multiple rules will need to be summarized concurrently. The approach of Wallis mentioned above (and outlined in the attached technical report) may provide a useful starting point for approaching this problem. We expect to begin to work on more comprehensive techniques for trend analysis and explanation during the coming year after we have completed the basic prototype system.

2 SCHEDULE FOR THE NEXT QUARTER (January 1, 1982 - March 31, 1982)

2.1 MULTIPLE REPRESENTATION SYSTEM (MRS)

2.1.1 Knowledge Representation

The major goal of our work on knowledge representation during the next quarter will be the development of MRS's meta-control capabilities. Mr. David Smith will perform much of this work. As described in the attached reports, it is possible to use MRS's language to express facts about MRS itself; the system can reason with these facts just as it reasons about external domains. Furthermore, MRS can examine and modify its own state. During the coming months we intend to exploit these capabilities to give MRS control of its search behavior, and, more specifically, the order in which it works on conjunctive goals or applies alternative methods.

While this research is proceeding, we intend to document and release the recently completed MRS compiler and truth maintenance facility. At present MRS reasons about each step before it performs it. While this is essential to MRS's control and modifiability, it is somewhat costly. The MRS compiler allows much of this reasoning to take place at the time of a modification and thereby saves repeated recomputation. The truth maintenance system automatically records justifications for all conclusions the systems draws, so that when a fact is removed from the data base, its consequences are automatically removed as well.

2.1.2 Problem Reformulation

Because Mr. Steve Tappel will be on leave of absence during the coming quarter, we do not anticipate significant further work on this problem until later in the year.

2.1.3 Causality, Design, and Mechanistic Reasoning

While SUBTLE is proving to be a powerful language for describing digital circuits, its vocabulary has not been developed beyond what is needed for that kind of application. For example, there are no terms for describing feedback or oscillation in analog circuits, and there are no terms for describing fluid flow, spatial organization, and the other topics relevant to non-electronic devices. During the coming quarter we hope to extend SUBTLE's vocabulary in these directions, largely through John Kunz's work on the encoding of renal physiology.

SCHEDULE FOR THE NEXT QUARTER (January 1, 1982 - March 31, 1982)

2.2 EXPERT EXPLANATIONS PROJECT (EXPEX)

2.2.1 Physiologic Modeling

In the next quarter Mr. Kunz's work on physiologic modeling will proceed to the development of an interactive program that responds to simple questions about basic mechanistic principles. The prototype system will focus on fluid balance in the human body and will be able to respond to questions such as:

- (a) "Is the patient oliguric? Define oliguria? What are its causes, consequences, and manifestations? How important are these consequences?"

In the subsequent months we will add the capability for the program to respond to questions such as the following:

- (b) "What is the effect of infusing hyperosmotic saline? How much should be used to return plasma volume to normal? How fast will normalization occur?"

The knowledge structures necessary to answer questions of type (a) have already been designed and largely implemented. Questions of type (b) will build on the other structures, incorporating the kinds of mathematical descriptions of human physiology necessary to predict the effects of quantitative deviations from normal.

2.2.2 Explanations from Causal Relationships

Although Dr. Wallis has left Stanford to begin an internship, his work and other recent research in explanation continues to intrigue us. His approach to customizing explanations for different classes of users, and the work of Weiner for summarizing multiple reasoning steps [3], are being studied by Dr. Glenn Rennels in the context of the EMYCIN system [2]. Expert systems built using EMYCIN generate explanations based almost entirely on the translation of single rules. During the next quarter we expect to have designed new modules for EMYCIN that will permit it to summarize chains of rules in a single step. A future goal is to allow EMYCIN to modify the content of an explanation in response to an internal model of the system's user.

2.2.3 Hypothesis Analysis for Decision Support

Over the next several months Mr. Cooper will be increasing the sophistication of the analysis system based on the knowledge base from INTERNIST/CADUCEUS. Key improvements will be to allow multiple disease hypotheses and hypotheses that contain causal links between diseases. The latter will be a first step toward introducing time into the analysis system.

SCHEDULE FOR THE NEXT QUARTER (January 1, 1982 - March 31, 1982)

The adequacy of the explanations generated by the system will be measured by comparing them to some of the cases being analyzed by physicians in the study being conducted by Randy Teach (see below). All 12 cases are in the domain of endocrinology, so our focus will also be on that domain (in particular, disorders of calcium balance). It is likely that our prototype analysis system will initially fall far short of the physician's performance. The comparative assessment will thus be used to refine our view of how to improve the system.

2.2.4 Psychological Studies of Explanation

The data collection phase of Mr. Teach's study is now underway and will be completed during the next quarter. All six subjects have already worked on several of the endocrinology problems, and the remaining cases will be undertaken in the next four to six weeks. We expect the subsequent data analysis to be completed during the second quarter of 1982.

2.2.5 Temporal Reasoning and Explanation

Until now our work on ONCOCIN has emphasized the development of the underlying knowledge structures (rules, control blocks, and a network of basic hierarchic concepts) upon which our research into temporal reasoning and explanation will be based. Intensive work on expanding the reasoning and explanation capabilities of the system will begin after the basic prototype is complete, probably in the second or third quarter of this year. During the coming quarter we will begin to work on polishing the explanations generated from the knowledge representation and reasoning approach already in use.

References

References

1. Shortliffe, E.H., Scott, A.C., Bischoff, M.B., Campbell, A.B., van Melle, W., and Jacobs, C.B. ONCOCIN: An expert system for oncology protocol management. Proceedings of 7th Intl. Joint Conf. on Artificial Intelligence, pp. 876-881, Vancouver, B.C., August 1981.
2. vanMelle, W. A Domain Independent System that Aids in Constructing Knowledge-Based Consultation Programs. Doctoral dissertation, Computer Science Department, Stanford University, June 1980.
3. Weiner, J.L. BLAH: A system which explains its reasoning. Artificial Intelligence 15:19-48 (1980).

Distribution Statement

Defense Documentation Center Cameron Station Alexandria, VA 22314	12 copies
Office of Naval Research Arlington, VA 22217	
Information Systems Program (437)	2 copies
Code 200	1 copy
Code 455	1 copy
Code 458	1 copy
Office of Naval Research Eastern/Central Regional Office Bldg. 114 Section D 666 Summer St. Boston, MA 02210	1 copy
Office of Naval Research Branch Office, Chicago 536 South Clark St. Chicago, IL 60605	1 copy
Office of Naval Research Western Regional Office 1030 East Green St. Pasadena, CA 91106	1 copy
Naval Research Laboratory Technical Information Division, Code 2627 Washington, DC 20375	6 copies
Dr. A. L. Slafkosky Scientific Advisor Commandant of the Marine Corps (RD-1) Washington, DC 20380	1 copy
Naval Ocean Systems Center Advanced Software Technology Division Code 5200 San Diego, CA 92152	1 copy
Mr. E. H. Gleissner Naval Ship Research & Development Center Computation and Mathematics Department Bethesda, MD 20084	1 copy
Capt. Grace M. Hopper (008) Naval Data Automation Command Washington Navy Yard Bldg. 166 Washington, DC 20374	1 copy
Captain R. Martin 3311 West Avenue Newport News, VA 23607	1 copy

ONR Resident Representative
Durand Aeronautics Bldg., Rm. 165
Stanford University
Stanford, California 94305

1 copy

Sponsored Projects Office
Encina Hall
Stanford

1 copy*

Edward H. Shortliffe, M.D., Ph.D.
Stanford Med Ctr TC-117

1 copy

Michael R. Genesereth, Ph.D.
Computer Science Department
Stanford University

1 copy

Betty Scott
Computer Science Department
Stanford University

1 copy

*Annual report only

Titles and Authors of Enclosed Reports

12/1981

1. Wallis, J.S. and Shortliffe, E.H. Explanatory power for medical expert systems: studies in the representation of causal relationships for clinical consultations.
2. Genesereth, M.R., Grinberg, M. and Lark, J. SUBTLE Manual.
3. Genesereth, M.R., Greiner, R. and Smith, D.E. MRS Manual.
4. Genesereth, M.R. The architecture of a multiple representation system.